

Programming Self-Organizing Systems with the Higher-Order Chemical Language

JEAN-PIERRE BANÂTRE¹, PASCAL FRADET², YANN RADENAC¹

¹ INRIA / IRISA, Campus de Beaulieu, 35042 Rennes Cedex, France

² INRIA Rhône-Alpes, 655 avenue de l'Europe, 38330 Montbonnot, France

`jbanatre@irisa.fr`, `Pascal.Fradet@inria.fr`, `yradenac@irisa.fr`

Received on 17th March 2006

In a chemical language, computation is viewed as abstract molecules reacting in an abstract chemical solution. Data can be seen as molecules and operations as chemical reactions: if some molecules satisfy a reaction condition, they are replaced by the result of the reaction. When no reaction is possible within the solution, a normal form is reached and the program terminates. In this article, we introduce HOCL, the Higher-Order Chemical Language, where reaction rules are also considered as molecules. We illustrate the application of HOCL to the specification of self-organizing systems. We describe two case studies: an autonomic mail system and the coordination of an image-processing pipeline on a grid.

Key words: Chemical programming, higher-order multiset rewriting, self-organization, autonomic systems, coordination

1 INTRODUCTION

The chemical reaction metaphor has been discussed in various occasions in the literature. This metaphor describes computation in terms of a chemical solution in which molecules (representing data) interact freely according to reaction rules. Chemical solutions are represented by multisets (data-structure that allows several occurrences of the same element). Computation proceeds

by rewritings which consume and produce new elements according to conditions and transformation rules.

To the best of our knowledge, the Gamma formalism was the first “chemical model of computation” proposed as early as in 1986 [5] and later extended in [6]. A Gamma program is a collection of reaction rules acting on a multiset of basic elements. A reaction rule is made of a condition and an action. Execution proceeds by replacing elements satisfying the reaction condition by the elements specified by the action. The result of a Gamma program is obtained when a stable (or inert) state is reached, that is to say, when no reaction can take place anymore. For example, the reaction

$$\text{max} = \mathbf{replace } x, y \text{ by } x \text{ if } x \geq y$$

computes the maximum element of a non empty set. The reaction replaces any couple of elements x and y such that the reaction condition ($x \geq y$) holds by x . This process goes on till a stable state is reached, that is to say, when only the maximum element remains. The reaction

$$\text{primes} = \mathbf{replace } x, y \text{ by } y \text{ if } x \text{ div } y$$

computes the prime numbers lower or equal to a given number N when applied to the multiset of all numbers between 2 and N ($x \text{ div } y$ is true if and only if x divides y).

Let us emphasize the conciseness and elegance of Gamma programs. Programs can be expressed without artificial sequentiality (i. e., unrelated to the logic of the program). If several disjoint tuples of elements satisfy the condition, the reactions can be performed in parallel. The interested reader may find in [6] a long series of examples (string processing problems, graph problems, geometry problems, etc.) illustrating the Gamma programming style.

Let us point out that we consider the chemical reaction only as a metaphor allowing a fresh look at programs and computation. In the same spirit, the chemical reaction model has been used as a source of inspiration for many different works [2]. However, there are models, like artificial chemistries [8], which are closer to real chemical systems and, sometimes, even use real chemistry to compute. In contrast, our research focuses on the design of expressive and unconventional programming languages.

The Higher-Order Chemical Language (HOCL) is a higher-order extension of Gamma: chemical programs (reactions) are first-class citizens. HOCL increases the expressive power of Gamma by allowing one to write reactions that consume or produce other reactions. This feature greatly facilitates the

expression of self-organizing systems. The main objective of this article is to illustrate these benefits through case studies. We first introduce HOCL informally using simple examples. Then, we present its use to the specification of two self-organizing systems: an autonomic mail system and the coordination of an image-processing pipeline on a grid.

2 THE HIGHER-ORDER CHEMICAL LANGUAGE

HOCL [4] is a higher-order extension of Gamma based on the γ -calculus [3]. Here, we present briefly and informally the features of HOCL used in the case studies of Sections 3 and 4. The interested reader will find a much more complete and formal presentation in [4].

In HOCL, programs, solutions, data and reactions are molecules. A program is a solution of atoms

$$\langle A_1, \dots, A_n \rangle$$

that is, a multiset of atoms built using the associative and commutative operator “;”. Associativity and commutativity formalize the Brownian motion of a chemical solution. They can always be used to reorganize molecules. Atoms are either basic constants (integers, booleans, etc.), sub-solutions ($\langle M \rangle$), pairs ($A_1:A_2$) or reaction rules. A reaction rule is written

replace-one P by M if C

where P is a pattern which matches the required atoms, C is the reaction condition and M the result of the reaction. For example,

$$\langle \langle \text{replace-one } x::Int \text{ by } 9 \text{ if } x \geq 10 \rangle, 4, 9, 15 \rangle \rightarrow \langle 4, 9, 9 \rangle$$

That reaction rule resembles a ceiling function. Its pattern $x::Int$ matches an integer, the condition imposes the integer to be greater than 10 and the action replaces it by 9. In the rest of this article, we omit types in patterns when there is no ambiguity. For example, it is clear that the previous pattern must select an integer since the condition is $x \geq 10$; the previous reaction can be written **replace-one** x by 9 if $x \geq 10$ instead.

Such reaction rules are said to be *one-shot* since they are consumed when they react. Reaction rules which remain after they react are called *n-shot*. Like in Gamma, they are denoted by **replace** P by M if C . The execution of a chemical program consists in performing reactions (non deterministically

and possibly in parallel) until the solution becomes inert i. e., no reaction can take place anymore. For example, the following program computes the prime numbers lower than 10 using a chemical version of the Eratosthenes' sieve:

$$\langle\langle \mathbf{replace} \ x, y \ \mathbf{by} \ x \ \mathbf{if} \ x \ \mathbf{div} \ y \rangle, 2, 3, 4, 5, 6, 7, 8, 9, 10 \rangle$$

The reaction removes any element y which can be divided by another one x . Initially several reactions are possible. For example, the couple $(2, 10)$ can be replaced by 2, the couple $(3, 9)$ by 3 or $(4, 8)$ by 4 etc. The solution becomes inert when the rule cannot react with any couple of integers in the solution, that is to say, when the solution contains only prime numbers. The result of the computation in our example is $\langle\langle \mathbf{replace} \ x, y \ \mathbf{by} \ x \ \mathbf{if} \ x \ \mathbf{div} \ y \rangle, 2, 3, 5, 7 \rangle$.

A molecule inside a solution cannot react with a molecule outside the solution (the construct $\langle . \rangle$ can be seen as a membrane). Reaction rules can access the contents of a sub-solution only if it is inert. This important restriction introduces some sequentiality in an otherwise highly parallel model: all reactions should be performed in a sub-solution before its content may be accessed or extracted. So, the pattern $\langle P \rangle$ matches only inert solutions whose (inert) content matches the pattern P .

Reactions can be named (or tagged) using the syntax $\mathbf{name} = \mathbf{replace} \dots$. Names are used to match and extract specific reactions using the same syntax ($\mathbf{name} = x$). We often use the \mathbf{let} operator to name reactions and assume that

$$\mathbf{let} \ \mathbf{name} = M \ \mathbf{in} \ N \stackrel{\text{def}}{=} N[(\mathbf{name} = M)/\mathbf{name}]$$

that is, the occurrences of \mathbf{name} in N are replaced by $\mathbf{name} = M$.

We also often make use of the pattern ω which can match any molecule even the "empty one". This pattern is very convenient to extract elements from a solution.

Using all these features, the previous example can be rewritten in order to remove the reaction at the end of the computation:

$$\begin{aligned} & \mathbf{let} \ \mathbf{sieve} = \mathbf{replace} \ x, y \ \mathbf{by} \ x \ \mathbf{if} \ x \ \mathbf{div} \ y \ \mathbf{in} \\ & \mathbf{let} \ \mathbf{clean} = \mathbf{replace-one} \langle \mathbf{sieve} = x, \omega \rangle \ \mathbf{by} \ \omega \ \mathbf{in} \\ & \quad \langle \mathbf{clean}, \langle \mathbf{sieve}, 2, 3, 4, 5, 6, 7, 8, 9, 10 \rangle \rangle \end{aligned}$$

The reduction proceeds as follows:

$$\begin{aligned} & \langle \mathbf{clean} = \dots, \langle \mathbf{sieve} = \dots, 2, 3, 4, 5, 6, 7, 8, 9, 10 \rangle \rangle \\ & \xrightarrow{*} \langle \mathbf{clean} = \dots, \langle \mathbf{sieve} = \dots, 2, 3, 5, 7 \rangle \rangle \\ & \rightarrow \langle 2, 3, 5, 7 \rangle \end{aligned}$$

The reaction rule clean cannot be applied until the sub-solution is inert. Only sieve can react until all primes are computed. Then, the one-shot rule clean extracts the prime numbers and suppresses the reaction rule sieve.

3 AUTONOMIC SYSTEMS IN HOCL

New advances in networking and computing technology has produced an explosive growth in networked applications and information services. These applications are more and more complex, heterogeneous and dynamic by essence. This combination of new parameters results in application development, configuration and management which break present computing paradigms. These applications should be able to manage themselves and react without external intervention. The goal of Autonomic Computing is to realize computer and software systems and applications that can manage themselves in accordance with high-level guidance from humans [11, 12].

The Chemical paradigm is well suited to express autonomic properties. Consider the general problem of a system whose state must satisfy a number of invariant properties but which is submitted to external and uncontrolled changes. This system must constantly re-organize itself to satisfy the properties. Typically in a chemical setting, an invariant property is represented by a n -shot reaction rule aiming at reaching a stable state where the property is satisfied. An autonomic program is made of a collection of such rules which are applied (without any external intervention) as soon as the solution is unstable again. This way of expressing self-management is very practical and will be used intensively.

To support our claims, we now specify an autonomic mail system within the higher-order chemical framework. The system consists of mail servers, each one dealing with a particular address domain, and clients sending their messages to their domain server. Servers forward messages addressed to other domains to the network. They also get messages addressed to their domain from the network and direct them to the appropriate clients.

General description: self-organization. The mail system (see Figure 1) is described by a solution that uses several molecules:

- Messages exchanged between clients are represented by basic molecules whose structure is left unspecified. We just assume that relevant information (such as sender's address, recipient's address, etc.) can be extracted using appropriate functions (such as *sender*, *recipient*, *body*, *senderDomain*, *recipientDomain*, etc.).

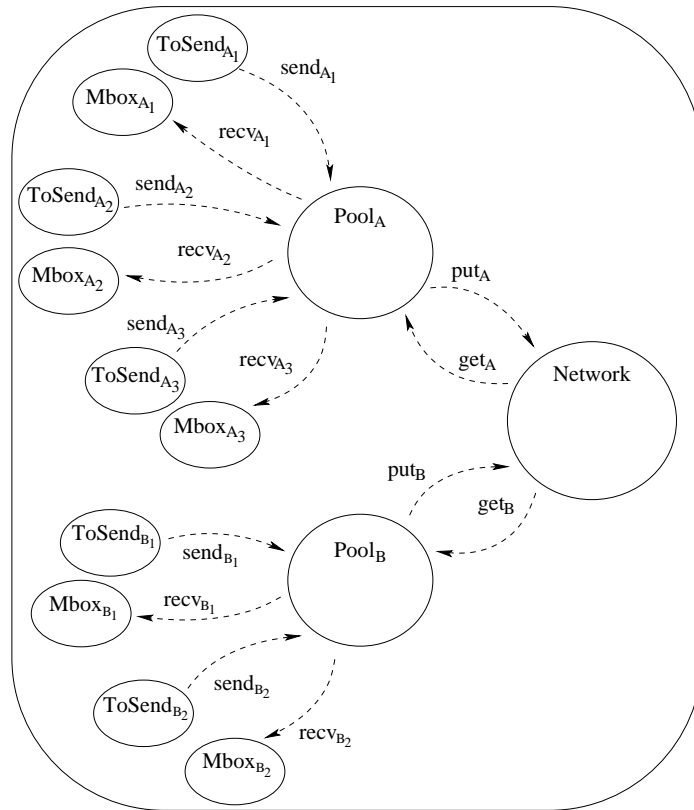


Figure 1
An autonomic mail system as a chemical solution.

- Solutions named $ToSend_{d_i}$ contain the messages to be sent by the client i of domain d .
- Solutions named $Mbox_{d_i}$ contain the messages received by the client i of domain d .
- Solutions named $Pool_d$ contain the messages that the server of domain d must take care of.
- The solution named Network represents the global network interconnecting domains.

- A client i in domain d is represented by two active molecules send_{d_i} and recv_{d_i} .
- A server of a domain d is represented by two active molecules put_d and get_d .

We do not specify within the system how new messages are produced by users. In any case, the mail system has no problem to deal with ToSend solutions changing dynamically. Clients send messages by adding them to the pool of messages of their domain. They receive messages from the pool of their domain and store them in their mailbox.

Messages stores are represented by sub-solutions. Movement of messages are made by reaction rules of the form:

$$\begin{array}{l} \mathbf{replace} \ A:\langle msg, \omega_A \rangle, \ B:\langle \omega_B \rangle \\ \mathbf{by} \ A:\langle \omega_A \rangle, \ B:\langle msg, \omega_B \rangle \\ \mathbf{if} \ Condition \end{array}$$

which moves the message msg from the store A to the store B if $Condition$ is satisfied.

Figure 2 gives the molecules that route messages in a self-organizing way. The send_{d_i} molecule sends messages of the client i (i.e., messages in the ToSend $_{d_i}$ solution) to the client's domain pool (i.e., the Pool $_d$ solution). The recv_{d_i} molecule places the messages addressed to client i (i.e., messages in the Pool $_d$ solution whose recipient is i) in the client's mailbox (i.e., the Mbox $_{d_i}$ solution). Servers forward messages from their pool to the network. They receive messages from the network and store them in their pool. The put_d molecule forwards only messages addressed to other domains than d . The molecule get_d extracts messages addressed to domain d from the network and places them in the pool of domain d . The system is a solution, named MailSystem, containing molecules representing clients, messages, pools, servers, mailboxes and the network. Figure 1 represents graphically the system with five clients grouped into two domains A and B .

Self-healing. We now assume that a server may crash. To prevent the mail service from being discontinued, we add an emergency server for each domain (see Figure 3 and 4). The emergency servers work with their own pool as usual but are active only when the corresponding main server has crashed. The modeling of a server crash can be done using the higher-order reaction rule crashServer_d : the active molecules representing a main server are replaced by molecules representing the corresponding emergency server. The

```

senddi = replace ToSenddi:⟨msg, ωt⟩, Poold:⟨ωp⟩
           by ToSenddi:⟨ωt⟩, Poold:⟨msg, ωp⟩

recvdi = replace Poold:⟨msg, ωp⟩, Mboxdi:⟨ωb⟩
           by Poold:⟨ωp⟩, Mboxdi:⟨msg, ωb⟩
           if recipient(msg) = i

putd = replace Poold:⟨msg, ωp⟩, Network:⟨ωn⟩
        by Poold:⟨ωp⟩, Network:⟨msg, ωn⟩
        if recipientDomain(msg) ≠ d

getd = replace Network:⟨msg, ωn⟩, Poold:⟨ωp⟩
        by Network:⟨ωn⟩, Poold:⟨msg, ωp⟩
        if recipientDomain(msg) = d

MailSystem:⟨
  sendA1, recvA1, ToSendA1:⟨...⟩, MboxA1:⟨...⟩,
  sendA2, recvA2, ToSendA2:⟨...⟩, MboxA2:⟨...⟩,
  sendA3, recvA3, ToSendA3:⟨...⟩, MboxA3:⟨...⟩,
  putA, getA, PoolA, Network, putB, getB, PoolB,
  sendB1, recvB1, ToSendB1:⟨...⟩, MboxB1:⟨...⟩,
  sendB2, recvB2, ToSendB2:⟨...⟩, MboxB2:⟨...⟩
⟩

```

Figure 2
Self-organizing routing molecules and the main solution.

boolean *failure* denotes a (potentially complex) failure detection mechanism. The inverse reaction repairServer_d represents the recovery of the server.

The two molecules Up_d and $(\text{DownIn}_d, \text{DownOut}_d)$ represent the state of the main server d in the solution. They are also active molecules in charge of transferring pending messages from Pool_d to $\text{Pool}_{d'}$; then, they may be forwarded by the emergency server.

The molecule DownOut_d transfers all messages bound to another domain than d from the main pool Pool_d to the emergency pool $\text{Pool}_{d'}$. The molecule DownIn_d transfers all messages bound to the domain d from the emergency pool $\text{Pool}_{d'}$ to the main pool Pool_d . After a transition from the down state to the up state, it may remain some messages in the emergency pools. So, the molecule Up_d brings back all the messages of emergency pool $\text{Pool}_{d'}$ into


```

crashServerd = replace putd, getd, Upd
                by putd', getd', DownInd, DownOutd
                if failure(d)
repairServerd = replace putd', getd', DownInd, DownOutd
                by putd, getd, Upd
                if recover(d)
DownOutd = replace Poold:⟨msg, ωp⟩, Poold':⟨ωn⟩
            by Poold:⟨ωp⟩, Poold':⟨msg, ωn⟩
            if domain(msg) ≠ d
DownInd = replace Poold:⟨ωp⟩, Poold':⟨msg, ωn⟩
            by Poold:⟨msg, ωp⟩, Poold':⟨ωn⟩
            if domain(msg) = d
Upd = replace Poold':⟨msg, ωp⟩, Poold:⟨ωn⟩
        by Poold':⟨ωp⟩, Poold:⟨msg, ωn⟩
MailSystem:⟨... , UpA, UpB, Pool'A, Pool'B,
            crashServerA, repairServerA,
            crashServerB, repairServerB⟩

```

Figure 3
Crash and self-healing molecules.

the main pool Pool_d to be then treated by the main server working again. In our example, self-healing can be implemented by two emergency servers A' and B' and boils down to adding the corresponding self-healing molecules to the solution.

Self-optimization. The emergency server can also be used to treat messages even if the main server is up. This improves efficiency by allowing parallelization. We can activate the emergency server when the main server is up with the molecules of Figure 5.

The role of the two molecules balance_d and $\text{balance}_{d'}$ is to perform dynamic load balancing between the two pools Pool_d and $\text{Pool}_{d'}$. Pools are balanced according to their number of messages (the function Card returns the cardinal of a molecule, i. e., the number of messages). We have to model the crash of the main server when the system is running in the optimized

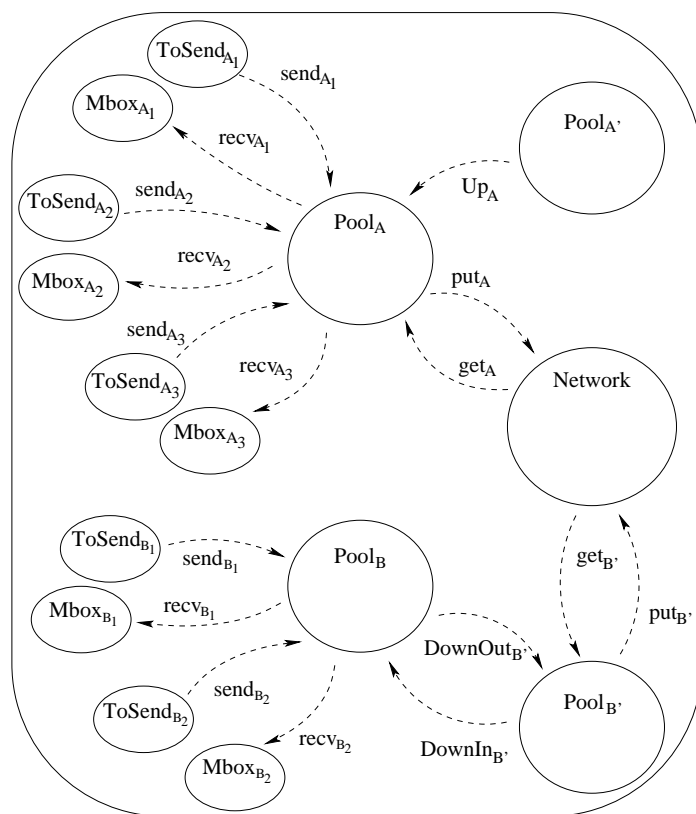


Figure 4
Highly-available mail system.

mode. The two molecules $balance_d$ and $balance_{d'}$ which characterize the optimized mode are removed and replaced by the molecules $DownIn_d$ and $DownOut_d$ which characterize the down state. After the main server is repaired the system may well trigger in the optimized mode again.

Adding, this self optimizing feature to our mail system boils down to adding these molecules to the main solution.

Self-protection. Self-protection can be decomposed in two phases: a detection phase and a reaction phase. The detection phase consists mainly in

```

optimized = replace Upd
           by putd', getd', balanced, balanced'
balanced = replace Poold:⟨msg, ωp⟩, Poold':⟨ωs⟩
           by Poold:⟨ωp⟩, Poold':⟨msg, ωs⟩
           if Card(ωp) > Card(ωs)
balanced' = replace Poold = ⟨ωp⟩, Poold':⟨msg, ωs⟩
            by Poold:⟨msg, ωp⟩, Poold':⟨ωs⟩
            if Card(ωp) < Card(ωs)
crashOptd = replace balanced, balanced'
            by DownInd, DownOutd
            if failure(d)

MailSystem:⟨... , optimizeA, optimizeB,
            crashOptA, crashOptB⟩

```

Figure 5
Self-optimizing molecules.

filtering data (pattern matching). The reaction phase consists in preventing offensive data from spreading and sometimes also in counter-attacking. This mechanism can easily be expressed with the condition-reaction scheme of the chemical paradigm. In our mail system, self-protection is simply implemented with active molecules of the following form:

$$\text{self-protect} = \mathbf{replace} \ x, \omega \ \mathbf{by} \ \omega \ \mathbf{if} \ filter(x)$$

If a molecule x is recognized as an offensive data by a filter function then it is suppressed. Variants of self-protect would consist in generating molecules to counter-attack or to send warnings.

Offensive data can take various forms such as spam, virus, etc. A protection against spam can be represented by the molecule:

$$\text{rmSpam} = \mathbf{replace} \ msg, \omega \ \mathbf{by} \ \omega \ \mathbf{if} \ isSpam(msg)$$

which is placed in a Pool_d solution. The contents of the pool can only be accessed when it is inert, that is when all spam messages have been suppressed by the active molecule rmSpam.

Self-configuration. We now consider adaptation and configuration issues that may arise with mobility. Assume that clients travel, move from personal computers to mobile phones, etc. Changes of environment suggest that clients should be able to migrate from a domain to another (closer or better suited to their new computing environment). We assume that the boolean $goTo_{d-e_i}$ signals to a client i in the domain d that it should go to the domain e . Such a migration can be described by the rule $migrate_{d-e_i}$ (see Figure 6).

$$\begin{aligned}
migrate_{d-e_i} &= \mathbf{replace} \text{ send}_{d_i}, \text{ recv}_{d_i}, \text{ Mbox}_{e_i}, \text{ ToSend}_{e_i}, \text{ Mbox}_{d_i}, \text{ ToSend}_{d_i} \\
&\quad \mathbf{by} \text{ send}_{e_i}, \text{ recv}_{e_i}, \text{ Mbox}_{d_i}, \text{ ToSend}_{d_i}, \langle \text{Fwd}:i:d:e \rangle \\
&\quad \mathbf{if} \text{ } goTo_{d-e_i} \\
forward_{d-e} &= \mathbf{replace} \text{ Pool}_d: \langle msg, \omega_p \rangle, \langle \text{Fwd}:i:d:e \rangle \\
&\quad \mathbf{by} \text{ Pool}_d: \langle newMsg(msg, e_i), \omega_p \rangle, \langle \text{Fwd}:i:d:e \rangle \\
&\quad \mathbf{if} \text{ } recipient(msg) = i \\
shortcut &= \mathbf{replace} \langle \text{Fwd}:i:a:b \rangle, \langle \text{Fwd}:i:b:c \rangle \\
&\quad \mathbf{by} \langle \text{Fwd}:i:a:c \rangle, \langle \text{Fwd}:i:b:c \rangle \\
shortcut' &= \mathbf{replace} \langle \text{Fwd}:i:d:d \rangle, \omega \mathbf{ by} \omega
\end{aligned}$$

Figure 6
Self-configuration molecules for migrations of users.

From now, the client will send and receive its messages from Pool_e . It may still have messages in its previous domain or some messages were sent before the migration and are in the network after the migration, or other clients may still send messages to its previous address. The migration places the tagged molecule $\langle \text{Fwd}:i:d:e \rangle$ in the solution in order to signal that messages to client i must be forwarded from domain d to domain e . The reaction rule $forward_{d-e}$ forwards messages between two domains d and e , where $newMsg(msg, r)$ builds a new message whose body is msg and recipient is r .

Notice that forward molecules accumulate when a client migrates several times (and remain even when the client is back to its original domain). We may prevent such forward chaining by using the molecules $shortcut$ and $shortcut'$.

Remarks. All these programs could (and should) be described generically (i. e., only once for all possible clients, source and destination domains). This would be easily done by tagging molecules (e. g., pools, clients, etc.). To simplify the presentation, we have described programs as if they were written for each client (server, domain, etc.). For the same reasons, we have also left tagging (e. g., of messages) implicit using extraction functions.

Our description should be regarded as a high-level parallel and modular specification. It allows to design and reason about autonomic systems at an appropriate level of abstraction. Reaction rules exhibit the essence of “autonomy” without going into useless details too early in the development process. Even if we have not designed a methodology for reasoning on higher-order chemical programs yet, basic reasoning can rely on the semantics and equivalence laws of HOCL defined in [4]. The interested reader may refer to [10] which presents a calculus of Gamma programs and its application to program reasoning and refinement.

Let us emphasize the conciseness and modularity of the resulting programs which rely essentially on the higher-order and chemical nature of HOCL. The self-healing, self-optimization and self-configuration programs are particularly illustrative of the elegance and power of the approach. They are described by independent collections of rules that are simply added to the system without requiring other changes. Note however that a direct implementation of this system is likely to be quite inefficient and further refinements are needed. This is another exciting research direction, not tackled here.

4 COORDINATION WITH HOCL

This section presents another example of a self-organizing system. The objective is to execute a program on several resources, more precisely, on a computational grid [9]. On each resource runs a process (or agent) that achieves a part of the computation of the program and coordinates with other processes (or agents) on other resources.

The basic application. Our application is an image-processing pipeline (IPP) composed of two filters F_1 and F_2 . For example, F_1 might reduce noise and a F_2 might increase contrast. The property to enforce is that filter F_1 must be applied before filter F_2 on each image. To ensure this, images are typed and filters are applied according to the type of images. Raw images

have type A, images returned by F_1 have type B and images returned by F_2 have type C. The HOCL program implementing that pipeline is:

```
let runF1 = replace x::A by (F1 x) in
let runF2 = replace x::B by (F2 x) in
⟨runF1, runF2, i1, i2, ...⟩
```

Initially, the solution has two reactions (runF₁ and runF₂) applying the filters and an arbitrary number of raw images i_1, i_2, \dots of type A.

Basic coordination on a grid. A grid is a network of resources which in principle, can be used transparently like a single machine to execute programs, store data, etc. A grid is a dynamic system: resources become overloaded, new ones become available, etc. So, before launching a program on a grid, it is in general impossible to know what resources will be allocated to the program. Resource allocation is determined dynamically according to their characteristics: CPU type, load state, computing power, memory size, connectivity, communication cost, etc.

A major challenge is to describe the coordination of the execution of programs on a grid in order to ensure some qualities of services (QoS) like efficiency, security, fault-tolerance, etc.

The chemical paradigm is well suited to the coordination of programs on a grid. Like in a chemical program where reactions occur in a parallel and chaotic way, communications (or coordinations) occur in a parallel and chaotic way between resources. Coordination between resources can be expressed by chemical reactions. Reaction rules specify different coordination that may occur in a solution of resources. For example, the system programmer may specify the dynamic placement or the migration of tasks using reaction rules. A load-balancing rule may specify that a task allocated to a resource should migrate to a less loaded resource.

The IPP on a grid. The image processing pipeline is described as a solution of resources each one being represented by a sub-solution. Each resource contains the programs that it runs. Reaction rules between sub-solutions represent the coordination between resources. Figure 7 gives a possible initial state of the system running the pipeline. The initial solution contains several resources R_i . A program is a pair of the form $name:\langle \dots \rangle$ when the solution contains tasks to execute and data to process. The reaction runFilter executes the tasks stored in the solution Tasks of the program. The resource R_1

```

let runFilter = replace r:⟨Pipe:⟨Tasks:⟨f:x, ωR⟩, ω1⟩, ω2⟩
                by r:⟨Pipe:⟨Tasks:⟨ωR⟩, (f x), ω1⟩, ω2⟩
    applyF1 = replace x::A, Tasks:⟨ωR⟩
                by Tasks:⟨F1:x, ωR⟩
    applyF2 = replace x::B, Tasks:⟨ωR⟩
                by Tasks:⟨F2:x, ωR⟩
    in
    ⟨runFilter,
     R1:⟨Pipe:⟨applyF1, applyF2, Tasks:⟨⟩, i1, i2, . . .⟩⟩, R2:⟨⟩, . . .⟩

```

Figure 7
A possible initial state of a grid runtime system for the pipeline.

contains the initial state of the Pipe program. First, the rules applyF₁ and applyF₂ find possible tasks (i. e., filter applications) and store them as pairs *filter:image* in the Tasks sub-solution. When the program becomes inert, the rule runFilter selects one of the stored tasks and runs it. The execution is finished when the top-level solution representing the system is inert. That is to say, when the solution R₁:⟨. . .⟩ is inert and its sub-solution Tasks is empty (no task remains).

This example represents resources, programs, tasks and data as molecules. However, it does not take profit from the grid (the application remains on a single resource). We now coordinate the execution of the IPP to ensure two QoS:

- *efficiency*, by parallelizing the execution according to the availability of resources,
- *security*, by encrypting/decrypting messages sent over untrusted communication links.

Parallelization. To ensure efficiency, the pipeline is spread on resources and executed in parallel. Different migration rules can be added to the system. Figure 8 gives such a rule (splitSPMD) and the rule transfer that moves messages (tasks and programs) between resources. In this part, messages

```

splitSPMD =
  let mergeSPMD =
    replace-one  $r_1:\langle\text{Pipe}:\langle\text{applyF}_1 = f, \text{applyF}_2 = g, \text{Tasks}:\langle\rangle, \omega_1\rangle, \omega\rangle,$ 
       $r_2:\langle\text{Pipe}:\langle\text{applyF}_1 = f', \text{applyF}_2 = g', \text{Tasks}:\langle\rangle, \omega_2\rangle, \omega'\rangle$ 
    by  $r_1:\langle\text{Pipe}:\langle\text{applyF}_1 = f, \text{applyF}_2 = g, \text{Tasks}:\langle\rangle, \omega_1\rangle, \omega\rangle,$ 
       $r_2:\langle\text{ToSend}:r_1:\langle\omega_2\rangle, \omega'\rangle$ 
  in
  replace  $r_1:\langle\text{Pipe}:\langle\text{applyF}_1 = f, \text{applyF}_2 = g, \text{Tasks}:\langle x\rangle, \omega_1\rangle, \omega'_1\rangle,$ 
     $r_2:\langle\omega_2\rangle$ 
  by  $r_1:\langle\text{Pipe}:\langle\text{applyF}_1 = f, \text{applyF}_2 = g, \text{Tasks}:\langle x_1\rangle, \omega_1\rangle,$ 
     $\text{ToSend}:r_2:\langle\text{Pipe}:\langle\text{applyF}_1 = f, \text{applyF}_2 = g, x_2\rangle\rangle, \omega'_1\rangle$ 
     $r_2:\langle\omega_2\rangle, \text{mergeSPMD}$ 
  if  $\text{Unbalanced}(r_1, r_2) \wedge (x_1, x_2) = \text{BalancedSplit}(r_1, r_2, x)$ 

transfer = replace  $r_1:\langle\text{ToSend}:r_2:m, \omega_1\rangle, r_2:\langle\omega_2\rangle$ 
  by  $r_1:\langle\omega_1\rangle, r_2:\langle m, \omega_2\rangle$ 

```

Figure 8
A possible distribution strategy of the pipeline.

could be left implicit but since they are central to the security QoS, we make them explicit using tuples of the form $\text{ToSend}:r:\text{message}$. The `splitSPMD` rule describes a possible distribution strategy of the pipeline. It splits the program in a Simple Program Multiple Data (SPMD) way. It is triggered by the predicate $\text{Unbalanced}(r_1, r_2)$ which is true when the load on r_1 is significantly higher than the load on r_2 . In this case, a part of (or maybe the whole) program that is running on r_1 is migrated on r_2 . This process is performed by the function $\text{BalancedSplit}(r_1, r_2, x)$ which takes a set of tasks (*filter:image*) and splits it into a set to be run on r_1 and a set to be run on r_2 . The program Pipe and the images to be transferred are transformed into messages of the form $\text{ToSend}:r:(\text{Pipe}:\langle\dots\rangle)$. The reaction rule `transfer` will move the message to the destination resource (r). When two resources have completed their tasks (the solution `Tasks` is empty and no message is

$$\begin{aligned}
\text{secureTransfer} &= \mathbf{replace} \ r_1:\langle \text{ToSend}:r_2:m, \omega_1 \rangle, \ r_2:\langle \omega_2 \rangle \\
&\quad \mathbf{by} \ r_1:\langle \omega_1 \rangle, \ r_2:\langle m, \omega_2 \rangle \\
&\quad \mathbf{if} \ \text{TrustedLink}(r_1, r_2) \\
&\quad\quad \vee (\neg \text{TrustedLink}(r_1, r_2) \wedge \text{Encrypted}(m)) \\
\text{encrypt} &= \mathbf{replace} \ r_1:\langle \text{ToSend}:r_2:m, \omega_1 \rangle \\
&\quad \mathbf{by} \ r_1:\langle \text{ToSend}:r_2:\text{Encrypt}(m), \omega_1 \rangle \\
&\quad \mathbf{if} \ \neg \text{TrustedLink}(r_1, r_2) \\
\text{decrypt} &= \mathbf{replace} \ r:\langle m, \omega \rangle \\
&\quad \mathbf{by} \ r:\langle \text{Decrypt}(m), \omega \rangle \\
&\quad \mathbf{if} \ \text{Encrypted}(m)
\end{aligned}$$

Figure 9
Rules for the security QoS.

pending), the reaction rule mergeSPMD merges their result (a resource sends its result to the other). This rule will gather all partial results onto a single resource.

Other migration and coordination rules may be defined and added to the system. For example, we could add a rule splitMPMD that splits the program in a Multiple Program Multiple Data (MPMD) way.

All these rule sets specify coordination rules and can be added to the single resource system described in Figure 7. Splits and merges may be applied in any order on any resources. The system self-organizes to keep the load balanced between resources. At the end, the top-level solution (the whole system) is inert and the result is available on a single resource.

Secure communications. Here, the required QoS is to ensure secure communications between resources. Either the communication takes place on a trusted link or otherwise messages are encrypted. We replace the rule transfer by the three rules secureTransfer, encrypt and decrypt (see Figure 9). The rules encrypt and decrypt encrypts and decrypts messages when needed. The rule secureTransfer moves the message only if the communication link is trusted or if it has been encrypted (by the rule encrypt). Like in the mail system, these reaction rules for ensuring security are modular. They can be

added to the system independently from the parallelization rules.

5 CONCLUSION

This article shows how the Chemical Programming paradigm can be used to describe self-organizing systems in a natural and elegant manner. In a first step, we introduced an expressive chemical language (called HOCL) whose higher-order properties allow the manipulation of programs. In HOCL, it becomes possible to write chemical reactions to create, delete or move around other chemical reactions. This feature is key in the description of self-organizing systems as shown in the two examples we have developed.

Self-organizing systems behave autonomously in order to maintain a pre-determined quality of service which may be violated in certain circumstances. Very often, such violations may be dealt with by applying local corrections according to pre-defined rules. These rules are easily expressed as HOCL reactions. The mail system as well as the coordination rules for grids are non trivial examples which illustrate self-organizing chemical computations.

To the authors' knowledge, there has been no other attempt to program self-organizing systems within the chemical paradigm framework. The closest contributions use rule-based languages such as Prolog to specify expert systems whose knowledge is represented as a collection of facts and (rewrite) rules. Some attempts have been reported concerning development of self-organizing applications on grids. The Project AutoMate [13] aims at defining self-management applications through the concept of autonomic elements which exhibit a self-organizing behavior. P-Grid [1] is a self-organizing peer-to-peer system. Peers run agents that store and retrieve data in the grid in a complete decentralized and autonomic way. In organic grids [7], the organization of computation is based on the autonomous scheduling of strongly mobile agents on a peer-to-peer network.

The contribution presented here is part of a more general research program on the use of chemical languages as coordination languages for the description of grid systems and applications. The basic challenge consists in showing that HOCL allows a clean and elegant expression of features such as program mobility, load balancing, crash recovery, etc. Basically, the overall system is expressed as a "soup" (represented by a multiset) of resources such as processors, storage, communication links, etc. whose organization is described by appropriate reaction rules. Further work consists in designing generic chemical operations to express more easily coordination (resource manipulation, task and data migration, etc.) in grids. Hopefully, programming a grid will be

as simple as specifying chemical solutions and reaction rules to describe how to distribute and execute them on a variety of resources.

REFERENCES

- [1] Kar Aberer. (2001). P-grid: A self-organizing access structure for p2p information systems. In *Proceedings of CoopIS 2001 conference*, number 2172 in LNCS. Springer-Verlag.
- [2] Jean-Pierre Banâtre, Pascal Fradet, and Daniel Le Métayer. (2001). Gamma and the chemical reaction model: Fifteen years after. In *Multiset Processing*, volume 2235 of LNCS, pages 17–44. Springer-Verlag.
- [3] Jean-Pierre Banâtre, Pascal Fradet, and Yann Radenac. (June 2004). Principles of chemical programming. In *Fifth International Workshop on Rule-Based Programming (RULE'04)*. Electronic Notes in Theoretical Computer Science 2005.
- [4] Jean-Pierre Banâtre, Pascal Fradet, and Yann Radenac. (November 2005). Generalized multisets for chemical programming. Research Report 5743, INRIA.
- [5] Jean-Pierre Banâtre and Daniel Le Métayer. (September 1986). A new computational model and its discipline of programming. Research Report 0566, INRIA.
- [6] Jean-Pierre Banâtre and Daniel Le Métayer. (January 1993). Programming by multiset transformation. *Communications of the ACM (CACM)*, 36(1):98–111.
- [7] Arjav J. Chakravarti, Gerald Baumgartner, and Mario Lauria. (2004). The organic grid: self-organizing computation on a peer-to-peer network. In *Proceedings of the Autonomic Computing Conference*, pages 96–103.
- [8] Peter Dittrich, Jens Ziegler, and Wolfgang Banzhaf. (2001). Artificial chemistries – a review. *Artificial Life*, 7(3):225–275.
- [9] Ian Foster and Carl Kesselman, editors. (1999). *The grid: blueprint for a new computing infrastructure*. Morgan Kaufmann Publishers Inc.
- [10] Chris Hankin, Daniel Le Métayer, and David Sands. (August 1992). A calculus of Gamma programs. In *Languages and Compilers for Parallel Computing, 5th International Workshop*, volume 757 of LNCS, pages 342–355. Springer-Verlag.
- [11] Jeffrey Kephart and David Chess. (January 2003). The vision of autonomic computing. *IEEE Computer*.
- [12] Manish Parashar and Salim Hariri. (2005). Autonomic computing: An overview. In Jean-Pierre Banâtre, Pascal Fradet, Jean-Louis Giavitto, and Olivier Michel, editors, *Unconventional Programming Paradigms*, volume 3566 of LNCS, pages 257–269. Springer.
- [13] Manish Parashar, Hua Liu, Zhen Li, Vincent Matossian, Cristina Schmidt, Guangsen Zhang, and Salim Hariri. (2006). AutoMate: Enabling autonomic grid applications. In *Cluster Computing: The Journal of Networks, Software Tools, and Applications, Special Issue on Autonomic Computing*, volume 9. Kluwer Academic.